

A NOTE ON CONVEX CHARACTERS, FIBONACCI NUMBERS AND EXPONENTIAL-TIME ALGORITHMS

STEVEN KELK AND GEORGIOS STAMOULIS

ABSTRACT. Phylogenetic trees are used to model evolution: leaves are labelled to represent contemporary species (“taxa”) and interior vertices represent extinct ancestors. Informally, convex characters are measurements on the contemporary species in which the subset of species (both contemporary and extinct) that share a given state, form a connected subtree. Given an unrooted, binary phylogenetic tree \mathcal{T} on a set of $n \geq 2$ taxa, a closed (but fairly opaque) expression for the number of convex characters on \mathcal{T} has been known since 1992, and this is independent of the exact topology of \mathcal{T} . In this note we prove that this number is actually equal to the $(2n-1)$ th Fibonacci number. Next, we define $g_k(\mathcal{T})$ to be the number of convex characters on \mathcal{T} in which each state appears on at least k taxa. We show that, somewhat curiously, $g_2(\mathcal{T})$ is also independent of the topology of \mathcal{T} , and is equal to the $(n-1)$ th Fibonacci number. As we demonstrate, this topological neutrality subsequently breaks down for $k \geq 3$. However, we show that for each fixed $k \geq 1$, $g_k(\mathcal{T})$ can be computed in $O(n)$ time and the set of characters thus counted can be efficiently listed and sampled. We use these insights to give a simple but effective exact algorithm for the NP-hard *maximum parsimony distance* problem that runs in time $\Theta(\phi^n \cdot n^2)$, where $\phi \approx 1.618\dots$ is the golden ratio, and an exact algorithm which computes the *tree bisection and reconnection distance* (equivalently, a *maximum agreement forest*) in time $\Theta(\phi^{2n} \cdot \text{poly}(n))$, where $\phi^2 \approx 2.619$.

1. INTRODUCTION

Phylogenetics is the science of accurately and efficiently inferring evolutionary trees given only information about contemporary species [12]. An important concept within phylogenetics is *convexity*. Essentially this captures the situation when, within a phylogenetic (i.e. evolutionary) tree, each biological state emerges exactly once: it should not emerge, die out, and then re-emerge. More concretely, given a phylogenetic tree and a set of states assigned to its leaves, can we assign states to the internal vertices of the tree such that each state forms a connected “island” within the tree? If this is possible, the assignment of states to the leaves is known as a *convex character*.

In this article we present a number of results concerning the enumeration of convex characters. In Section 2 we give formal definitions and describe relevant earlier work. In Section 3 we start by showing that an earlier result counting convex characters can be simplified to a term of the Fibonacci sequence. We then seek to count convex characters with the added restriction that each state should occur on at least k leaves, proving the somewhat surprising result that (as for $k = 1$) tree topology is irrelevant for $k = 2$, and that a formulation in terms of Fibonacci numbers is again possible. We give an explicit example showing that for $k \geq 3$ the topological neutrality breaks down. In Section 4 we show that for all k the size of the space can be counted in polynomial time and space using dynamic programming, which also permits listing and sampling uniformly at random, noting also that non-isomorphic trees can have exactly the same vector of space sizes (for $k = 1, 2, \dots$). In Section 5 we give a number of algorithmic applications for NP-hard problems arising in phylogenetics that seek to quantify the dissimilarity of two phylogenetic trees. Finally, in Section 6 we briefly discuss a number of open problems arising from this work. The software associated with this article has been made publicly available.

2. PRELIMINARIES

For general background on mathematical phylogenetics we refer to [12, 5]. An *unrooted binary phylogenetic X -tree* is an undirected tree $\mathcal{T} = (V(\mathcal{T}), E(\mathcal{T}))$ where every internal vertex has degree 3 and whose leaves are bijectively labelled by a set X , where X is often called the set of

taxa (representing the contemporary species). We use n to denote $|X|$ and often simply write *tree* when this is clear from the context.

A *character* f on X is a surjective function $f : X \rightarrow \mathcal{C}$ for some set \mathcal{C} of *states* (where a state represents some characteristic of the species e.g. number of legs). We say that f is an r -state character if $|\mathcal{C}| = r$. Each character naturally induces a partition of X and here we regard two characters as being equivalent if they both induce the same partition of X . An *extension* of a character f to $V(\mathcal{T})$ is a function $h : V(\mathcal{T}) \rightarrow \mathcal{C}$ such that $h(x) = f(x)$ for all x in X . For such an extension h of f , we denote by $l_h(\mathcal{T})$ the number of edges $e = \{u, v\}$ such that $h(u) \neq h(v)$. The *parsimony score* of a character f on \mathcal{T} , denoted by $l_f(\mathcal{T})$, is obtained by minimizing $l_h(\mathcal{T})$ over all possible extensions h of f . We say that a character $f : X \rightarrow \mathcal{C}$ is *convex on \mathcal{T}* if $l_f(\mathcal{T}) = |\mathcal{C}| - 1$. Equivalently: a character $f : X \rightarrow \mathcal{C}$ is convex on \mathcal{T} if there exists an extension h of f such that, for each state $c \in \mathcal{C}$, the vertices of \mathcal{T} that are allocated state c (by h) form a connected subtree of \mathcal{T} . We call such an extension h a *convex extension* of f . See Figure 1 for an example. The convexity of a character can be tested in polynomial [7, 8] (in fact, linear [2]) time.




FIGURE 1. For the given tree \mathcal{T} (on 7 taxa) there are 233 convex characters in total, but only 8 in which each state appears on at least 2 taxa, and these are shown above. 1 character uses exactly 1 state ($abcdefg$), so $g_2(\mathcal{T}, 1) = 1$, 4 characters use 2 states ($ab|cdefg$, $abc|defg$, $abcd|efg$ and $abcde|fg$), so $g_2(\mathcal{T}, 2) = 4$ and 3 characters use 3 states ($ab|cd|efg$, $abc|de|fg$ and $ab|cde|fg$), so $g_2(\mathcal{T}, 3) = 3$. For each character we have shown an extension verifying that the subtree induced by each state is connected i.e. that the character is convex.

We write $g(\mathcal{T}, r)$ to denote the number of r -state convex characters on \mathcal{T} and $g_i(\mathcal{T}, r)$ ($i \geq 1$) to denote the number of those characters that have the additional property that each state used by the character appears on at least i taxa. It follows from the definition of character that $g_1(\mathcal{T}, r) = g(\mathcal{T}, r)$. We define:

$$g_i(\mathcal{T}) = \sum_{r=1}^n g_i(\mathcal{T}, r).$$

The value $g_1(\mathcal{T})$ is therefore equal to the total number of convex characters on \mathcal{T} . For the tree shown in Figure 1, $g(\mathcal{T}) = g_1(\mathcal{T}) = 233$ and $g_2(\mathcal{T}) = 8$. We adopt the standard convention that the binomial coefficient $\binom{n}{k}$ evaluates to 1 if $k = 0$, and 0 if $n < k$ and $k > 0$. In [13] it is proven that, for $n, r \geq 1$,

$$g(\mathcal{T}, r) = g_1(\mathcal{T}, r) = \binom{2n - r - 1}{r - 1}.$$

Hence,

$$(1) \quad g_1(\mathcal{T}) = \sum_{r=1}^n \binom{2n-r-1}{r-1}.$$

As observed in [13], the expression for $g_1(\mathcal{T}, r)$ (somewhat surprisingly) does not depend on the topology of \mathcal{T} , only on the number of taxa n . Hence we can write $g(n)$, $g_1(n)$ and $g_1(n, r)$ without ambiguity.

3. FIBONACCI NUMBERS AND CONVEX CHARACTERS

Theorem 1. *The value $g_2(\mathcal{T}, r)$ only depends on n (i.e. the topology of \mathcal{T} is not relevant) and for $n \geq 2, r \geq 1$ is given by the expression*

$$g_2(\mathcal{T}, r) = \binom{n-r-1}{r-1}.$$

Proof. We prove this by induction on n . For the base case note that for $n \in \{2, 3, 4\}$ there is only one binary tree topology (up to relabelling of taxa) possible on n taxa and that the expression correctly evaluates to 1 when $r = 1$ and, when $r \geq 2$, evaluates to 0 in all cases except $n = 4, r = 2$ when it correctly evaluates to 1.

Consider then $n \geq 5$, and let r be any value $2 \leq r \leq \lfloor n/2 \rfloor$. (For $r = 1$ the expression correctly evaluates to 1, and for $r > \lfloor n/2 \rfloor$ the expression correctly evaluates to 0). Every tree on 4 or more taxa contains at least one *cherry*: two taxa x, y that have a common parent u where the third neighbour of u is an interior vertex. Fix such a cherry. (A similar technique is used in [13] and [14]). Observe that any convex character f on \mathcal{T} with the property that each state appears at least twice, has $f(x) = f(y)$. This follows from the connected-subtree definition of convexity. Now, let $\mathcal{T}' = \mathcal{T}|(X \setminus \{x\})$ and let $\mathcal{T}'' = \mathcal{T}|(X \setminus \{x, y\})$, where $\mathcal{T}|X'$ denotes the tree (on the set of taxa X') obtained from \mathcal{T} by taking the minimum subtree connecting the elements of X' and then suppressing vertices of degree 2.

There are two cases to distinguish. The first case is when the state $f(x) = f(y)$ does not appear on any other taxa. There are $g_2(\mathcal{T}'', r-1)$ such characters. The second case is when $f(x) = f(y)$ does appear on at least one other taxon. There are $g_2(\mathcal{T}', r)$ such characters. Hence,

$$(2) \quad g_2(\mathcal{T}, r) = g_2(\mathcal{T}'', r-1) + g_2(\mathcal{T}', r).$$

By the inductive hypothesis we have

$$\begin{aligned} g_2(\mathcal{T}, r) &= \binom{(n-2)-(r-1)-1}{(r-1)-1} + \binom{(n-1)-r-1}{r-1} \\ &= \binom{n-r-2}{r-2} + \binom{n-r-2}{r-1} \\ &= \binom{n-r-1}{r-1}. \end{aligned}$$

The last equality follows from the well-known identity known as *Pascal's Rule* i.e. $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$, which holds for $1 \leq k \leq n$. This completes the proof. \square

Consequently, the total number of convex characters on a tree \mathcal{T} with each state appearing at least twice is independent of the topology of \mathcal{T} . Specifically, we have:

$$(3) \quad g_2(\mathcal{T}) = g_2(n) = \sum_{r=1}^n \binom{n-r-1}{r-1} = \sum_{r=1}^{\lfloor n/2 \rfloor} \binom{n-r-1}{r-1}$$

Corollary 2. *For even n , $g_2(n) = g_1(n/2)$.*

Proof. This is immediate by observing that Equation (3) can be obtained by substituting $n/2$ for n in Equation (1). \square

Let $F(n)$ denote the n th Fibonacci number. That is, $F(0) = 0$, $F(1) = 1$ and for $n \geq 2$, $F(n) = F(n-1) + F(n-2)$. For comprehensive background on Fibonacci numbers see [11].

Theorem 3. For $n \geq 2$, $g_2(n) = F(n-1)$ and $g_1(n) = F(2n-1)$.

Proof. The following identity is classical¹ ($n \geq 0$):

$$F(n+1) = \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n-k}{k}.$$

If we index k from 1 rather than 0 we obtain,

$$F(n+1) = \sum_{k=1}^{\lfloor n/2 \rfloor + 1} \binom{n-(k-1)}{k-1}.$$

Now, if we replace n with $n-2$:

$$\begin{aligned} F(n-1) &= \sum_{k=1}^{\lfloor (n-2)/2 \rfloor + 1} \binom{n-k-1}{k-1} \\ &= \sum_{k=1}^{\lfloor n/2 \rfloor} \binom{n-k-1}{k-1} \\ &= g_2(n). \end{aligned}$$

The expression for $g_1(n)$ is then obtained by applying Corollary 2. □

The question arises whether the values $g_3(\mathcal{T}, r)$ and/or $g_3(\mathcal{T})$ share the topological neutrality of their g_1 and g_2 counterparts. This is not the case: see Figure 2. Here $g_3(\mathcal{T}_1) = 5$, because $g_3(\mathcal{T}_1, 1) = 1$, $g_3(\mathcal{T}_1, 2) = 3$, $g_3(\mathcal{T}_1, 3) = 1$ and $g_3(\mathcal{T}_1, r) = 0$ (for $r > 3$). However, $g_3(\mathcal{T}_2) = 6$, because $g_3(\mathcal{T}_2, 1) = 1$, $g_3(\mathcal{T}_2, 2) = 4$, $g_3(\mathcal{T}_2, 3) = 1$ and $g_3(\mathcal{T}_2, r) = 0$ (for $r > 3$).




FIGURE 2. The number of characters that are convex on \mathcal{T}_1 with each state appearing on at least 3 taxa, is not the same as the corresponding number for \mathcal{T}_2 i.e. $g_3(\mathcal{T}_1) \neq g_3(\mathcal{T}_2)$. This is because $g_3(\mathcal{T}_1, 2) = 3$ (the characters $abc|defghi$, $abcdef|ghi$ and $abceghi|def$) while $g_3(\mathcal{T}_2, 2) = 4$ (the characters $abc|defghi$, $abcd|efghi$, $abcde|fgh$ and $abcdef|ghi$). Hence topology does play a role here, contrasting with the situation for g_1 and g_2 .

¹This is usually attributed to Édouard Lucas (1876). By applying Pascal's Rule and some algebraic manipulation it can be proven by induction.

4. COMPUTING $g_k(\mathcal{T})$ (AND LISTING ITS ELEMENTS) WITH DYNAMIC PROGRAMMING

The results from the previous section give rise to a number of questions. Can we compute $g_k(\mathcal{T})$ in polynomial time, for $k \geq 3$? Also, if we want to explicitly list all the elements counted by $g_k(\mathcal{T})$ ($k \geq 1$), is it possible to achieve this in some reasonable total running time e.g. $O(g_k(\mathcal{T}) \cdot \text{poly}(n))$? We show that, for all $k \geq 1$, the answer to both questions is *yes*. Specifically, we show how to compute $g_k(\mathcal{T})$ using dynamic programming, for each $k \geq 1$. The combinatorial recurrence within the dynamic programming will also allow us to derive a polynomial-time computable bijection from $\{1, 2, \dots, g_k(\mathcal{T})\}$ to the characters counted by $g_k(\mathcal{T})$. Using this bijection it is then straightforward to list (or sample) these elements. Note that this is also an advance for $g_1(\mathcal{T})$, since the recurrence used in [13] to derive $g_1(n)$ is based on inclusion-exclusion: it is not obvious how to transform it into a bijection.

We begin by rooting \mathcal{T} by subdividing an arbitrary edge with a new vertex and (implicitly) directing all edges away from this new vertex. The new vertex becomes the root of the tree. Note that this rooting operation has no impact on the convexity of characters, and the location of the root is irrelevant; it is simply a convenience which ensures that the term “child” is well-defined. The dynamic programming works bottom-up, from the leaves towards the root.

Here it is helpful to represent a character f on X as a set of non-overlapping, non-empty subsets $\{B_1, \dots, B_t\}$ that partition X , where each B_i corresponds to a state. We also need some new definitions. A character f is *valid* for $g_k(\mathcal{T})$ if

- f is convex on \mathcal{T} , and
- for each $B_i \in f$, $|B_i| \geq k$.

Consider an ordered pair (f, B) where f is a character on X and $B \in f$. We call such a pair a *character-root* pair of \mathcal{T} if

- f is convex on \mathcal{T} , and
- there exists a convex extension of f in which the root of \mathcal{T} is assigned state B .

Equality between character-root pairs is defined strictly i.e. $(f, B) = (f', B')$ if and only if $f = f'$ and $B = B'$. We say that a character-root pair (f, B) of \mathcal{T} is *semi-valid* for $g_k(\mathcal{T})$ if, for each $B_i \in f$, $B_i \neq B \Rightarrow |B_i| \geq k$. Note that if a character-root pair (f, B) is semi-valid for $g_k(\mathcal{T})$, then f is valid for $g_k(\mathcal{T})$ if and only if $|B| \geq k$.

At each vertex u of the tree we will compute and store the following $k + 1$ values, where \mathcal{T}_u is simply the subtree rooted at u :

- $g_k(\mathcal{T}_u)$,
- for each $1 \leq m \leq k - 1$ the number $h(\mathcal{T}_u, m)$ which is defined as the number of character-root state pairs (f, B) of \mathcal{T}_u such that both the following conditions hold: (f, B) is semi-valid for $g_k(\mathcal{T}_u)$ and $|B| = m$. We also store $h(\mathcal{T}_u, k)$ which is defined slightly differently: we replace the term $|B| = m$ with $|B| \geq k$.

If u is a taxon, then:

- $g_k(\mathcal{T}_u)$ is equal to 1 if $k = 1$, and 0 if $k > 1$,
- $h(\mathcal{T}_u, m)$ is equal to 1 if $m = 1$ and 0 if $m > 1$.

We show how to compute these values recursively, assuming the corresponding values have already been computed for \mathcal{T}_l (the subtree rooted at the left child of u) and \mathcal{T}_r (the subtree rooted at the right child of u). First,

$$(4) \quad g_k(\mathcal{T}_u) = g_k(\mathcal{T}_l)g_k(\mathcal{T}_r) + \sum_{\substack{1 \leq i, j \leq k \\ i+j \geq k}} h(\mathcal{T}_l, i)h(\mathcal{T}_r, j).$$

The idea behind this recurrence is that characters counted by $g_k(\mathcal{T}_u)$ can be created in two ways: (i) by taking the union of a character from the left subtree with a character from the right subtree, and (ii) by taking a character-root pair (f, B) from the left subtree, a character-root pair (f', B') from the right subtree, and then merging the root states to yield a character $(f \setminus B) \cup (f' \setminus B') \cup (B \cup B')$. Characters from the subtrees can only be used for (i) if they are already valid (with respect to g_k) in their subtrees. The characters f and f' that are used for (ii) might not be valid with respect

to their subtrees, but we do require that they can be combined to obtain a character that *is* valid for $g_k(\mathcal{T}_u)$. This is possible if and only if (f, B) and (f', B') are semi-valid for their respective subtrees and the sum of the cardinalities of B and B' is at least k . See Figure 3 for an example.




FIGURE 3. The character $abc|d$ is not valid for $g_3(\mathcal{T}_l)$, and $efg|hi$ is not valid for $g_3(\mathcal{T}_r)$, but the character $abc|dhi|efg$ is valid for $g_3(\mathcal{T}_u)$. The validity is obtained by allowing d and hi to merge, which was possible because they could both “reach” the roots of their respective subtrees.

Second, for $1 \leq m \leq k - 1$ we have,

$$(5) \quad h(\mathcal{T}_u, m) = g_k(\mathcal{T}_l)h(\mathcal{T}_r, m) + h(\mathcal{T}_l, m)g_k(\mathcal{T}_r) + \sum_{\substack{1 \leq i, j \leq k-2 \\ i+j=m}} h(\mathcal{T}_l, i)h(\mathcal{T}_r, j).$$

Note that here $m \leq k - 1$, which means that the semi-valid character-root pairs (f, B) counted by this recurrence are such that f is *not* valid for $g_k(\mathcal{T}_u)$. The first two terms of the recurrence concern the situation analogous to (i) above. Specifically, in this case we assume that no states are merged, so a new semi-valid character-root pair can be created for \mathcal{T}_u if and only if it is constructed from the combination of a valid character from one subtree, with a semi-valid character-root pair from the other. The summation term corresponds to (ii). That is, we only count combinations of character-root pairs from the two subtrees if the cardinality of the merged state is exactly m .

Finally we have,

$$(6) \quad h(\mathcal{T}_u, k) = g_k(\mathcal{T}_l)h(\mathcal{T}_r, k) + h(\mathcal{T}_l, k)g_k(\mathcal{T}_r) + \sum_{\substack{1 \leq i, j \leq k \\ i+j \geq k}} h(\mathcal{T}_l, i)h(\mathcal{T}_r, j).$$

This final recurrence is semantically very similar to the previous one. The main difference is that it counts all semi-valid character-root pairs (f, B) for $g_k(\mathcal{T}_u)$ such that f *is* valid for $g_k(\mathcal{T}_u)$.

For a given vertex u , Equation (4) can be computed in $O(k^2)$ time, assuming the values for \mathcal{T}_l and \mathcal{T}_r have already been computed earlier. The same time bound holds for Equations (5) (for a specific $1 \leq m \leq k - 1$) and (6). Equation (5) has to be computed for each m , yielding a naïve running time bound of $O(k^3)$ (per vertex u), but this can easily be improved to $O(k^2)$ by observing that a single $1 \leq i, j \leq k$ sweep over the $h(\mathcal{T}_l, i)$ and $h(\mathcal{T}_r, j)$ values can be recycled for computation of all the different $h(\mathcal{T}_u, m)$ values. There are $2(n - 1)$ vertices in the tree. This yields the following theorem.

Theorem 4. *Let \mathcal{T} be an unrooted binary tree on n taxa. For each $1 \leq k \leq n$, $g_k(\mathcal{T})$ can be computed in $O(k^2 \cdot n)$ time and $O(k \cdot n)$ space.*

Corollary 5. *Let \mathcal{T} be an unrooted binary tree on n taxa. For each $1 \leq k \leq n$, all the characters that are counted by $g_k(\mathcal{T})$ can be generated in $O(g_k(\mathcal{T}) \cdot k^2 \cdot n)$ total time, and a character counted by $g_k(\mathcal{T})$ can be sampled uniformly at random in $O(k^2 \cdot n)$ time and $O(k \cdot n)$ space.*

Proof. Critically, there is no inclusion-exclusion involved Equations (4), (5) and (6). This allows us to impose a canonical ordering on the characters (and character-root pairs) counted by these

equations. For example, within Equation (4) we can choose to place the type-(i) characters earlier in the ordering than the type-(ii) characters. Within the $g_k(\mathcal{T}_l)g_k(\mathcal{T}_r)$ type-(i) characters we can refine the order as follows: the *first* character from the left subtree combined in turn each of the $g_k(\mathcal{T}_r)$ characters from the right subtree, then the *second* character from the left subtree combined in turn with each of the $g_k(\mathcal{T}_r)$ characters from the right subtree, and so on. Once a canonical ordering has been chosen and the dynamic programming has been completed, we can start at the root of \mathcal{T} and (using the g and h values computed at all vertices of the tree) recursively backtrack to generate the uniquely defined i th character. Hence, we obtain a bijection from $\{1, 2, \dots, g_k(\mathcal{T})\}$ to the characters counted by $g_k(\mathcal{T})$. The time and space requirements for backtracking through the tree (i.e. evaluating the bijection for a given element of $\{1, 2, \dots, g_k(\mathcal{T})\}$) are dominated by the time and space requirements of executing the original dynamic program, which are $O(k^2 \cdot n)$ and $O(k \cdot n)$ respectively. This bijection can then be used to list all the characters counted by $g_k(\mathcal{T})$ or to sample uniformly at random from this space. \square

We have implemented the dynamic programming (and the corresponding algorithms for listing and sampling) in Java and this can be downloaded from <http://skelk.sdf-eu.org/convexcount>.

Finally within the section, for an unrooted binary tree \mathcal{T} on n leaves, we define the g -spectrum as simply the vector $(g_1(\mathcal{T}), g_2(\mathcal{T}), \dots, g_n(\mathcal{T}))$. It is natural to ask whether two trees on n taxa have the same g -spectrum if and only if they are isomorphic (see e.g. [3] for related discussions of “identifiability”). Using the code above we have verified that, while this claim is true for $n \leq 10$ leaves (see the software website for a proof), a counter-example exists for $n = 11$, see Figure 4.




FIGURE 4. These two trees on 11 leaves are non-isomorphic but have the same g -spectrum: $(10946, 55, 8, 3, 2, 1, 1, 1, 1, 1, 1)$.

5. ALGORITHMIC APPLICATIONS

One of the advantages of expressing $g_1(n)$ and $g_2(n)$ as Fibonacci numbers is that it allows us to give tight bounds on their rate of growth. This can be particularly useful when bounding the running time of algorithms. Consider the following classical, closed-form expression for the Fibonacci numbers ($n \geq 0$), where $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618\dots$ is the *golden ratio*:

$$F(n) = \left\lfloor \frac{\phi^n}{\sqrt{5}} + \frac{1}{2} \right\rfloor.$$

(It is obtained from Binet's Formula $F(n) = \frac{\phi^n - (-\phi)^{-n}}{\sqrt{5}}$ (1843) by observing that the term $(-\phi)^{-n}$ is vanishing.) Combining with Theorem 3 we obtain

$$g_1(n) = \left\lfloor \frac{\phi^{2n-1}}{\sqrt{5}} + \frac{1}{2} \right\rfloor,$$

$$g_2(n) = \left\lfloor \frac{\phi^{n-1}}{\sqrt{5}} + \frac{1}{2} \right\rfloor.$$

Using asymptotic notation, it is clear that there are $\Theta(\phi^{2n})$ convex characters and $\Theta(\phi^n)$ convex characters in which each state occurs on at least two taxa. We give two examples of how these insights yield non-trivial exponential-time algorithms for two NP-hard problems arising in phylogenetics.

5.1. Computation of maximum parsimony distance. Let $\mathcal{T}_1, \mathcal{T}_2$ be two unrooted binary trees on the same set of taxa X . The metric $d_{MP}(\mathcal{T}_1, \mathcal{T}_2)$ (the *maximum parsimony distance* of \mathcal{T}_1 and \mathcal{T}_2) is defined as follows, where f ranges over *all* characters on X and $l_f(\cdot)$ is as defined in Section 2:

$$d_{MP}(\mathcal{T}_1, \mathcal{T}_2) = \max_f |l_f(\mathcal{T}_1) - l_f(\mathcal{T}_2)|$$

It is NP-hard to compute d_{MP} [9, 6]. It can be used to quantify the dissimilarity of two phylogenetic trees and is a lower bound on the similarly NP-hard *tree bisection and reconnection* (TBR) distance, denoted d_{TBR} [1].

Theorem 6. *Given two unrooted binary trees $\mathcal{T}_1, \mathcal{T}_2$ on the same set of taxa X , where $|X| = n$, $d_{MP}(\mathcal{T}_1, \mathcal{T}_2)$ can be computed in time $\Theta(\phi^n \cdot n^2)$, where $\phi \approx 1.618\dots$ is the golden ratio.*

Proof. In [9, 6] it is proven that the optimum is achieved by some character f that is convex on \mathcal{T}_1 or \mathcal{T}_2 and where each state in the character occurs on at least two taxa. Hence simply looping through all the characters counted by $g_2(\mathcal{T}_1)$ and, separately, all the characters counted by $g_2(\mathcal{T}_2)$ is sufficient to locate an optimal character. Note that $l_f(\cdot)$ can be computed in $O(n^2)$ time using Fitch's algorithm² [7] or dynamic programming. Hence, scoring each character f can easily be performed in quadratic time. The result then follows by leveraging Corollary 5. \square

We have implemented the d_{MP} algorithm in Java and for an exponential-time algorithm the results are encouraging; the code is freely available at <http://skelk.sdf-eu.org/convexmpdist>. On a single 32-bit 1.66GHz Intel Atom (N450) processor the algorithm terminates for $n = 20, 25, 30$ in less than 1 second, 3 seconds and 51 seconds respectively. On a more powerful 64-bit 3.10GHz machine the previously fastest algorithm, the Integer Linear Programming (ILP) approach described in [9], took 70 seconds to terminate on 12 taxa, and stalled completely on trees with more than 16 taxa, even using state-of-the-art ILP software. The enhanced range of our software has been recently used in experiments to verify that d_{MP} is often a very good lower bound on d_{MP} [10].

5.2. Computation of TBR distance and maximum agreement forests. Finally, we note that the results in this article also give an easy (although, in some cases, somewhat crude) upper bound on the number of *agreement forests* of two unrooted binary trees $\mathcal{T}_1, \mathcal{T}_2$ on n taxa. Recall that, for an unrooted binary phylogenetic tree \mathcal{T} on X and $X' \subseteq X$, $\mathcal{T}|X'$ is defined to be the unrooted binary phylogenetic tree on X' obtained by taking the minimal subtree of \mathcal{T} that connects X' , and suppressing vertices of degree 2. An agreement forest is a partition of X into non-empty subsets X_1, \dots, X_k such that (i) within \mathcal{T}_1 (respectively, \mathcal{T}_2) the minimal connecting subtrees induced by the X_i are vertex-disjoint and (ii) for each X_i , $\mathcal{T}_1|X_i = \mathcal{T}_2|X_i$ (where here equality explicitly takes the taxa into account). See [1] and recent articles such as [4] for further background on agreement forests. A *maximum agreement forest* is an agreement forest with a

²Fitch's algorithm has running time $O(ns)$ where n is the number of taxa and s is the number of states in the character. In our context s can rise to $O(n)$.

minimum number of components, and this minimum is denoted d_{MAF} . Note that due to part (i) of the definition every agreement forest induces a character that is convex on both \mathcal{T}_1 and \mathcal{T}_2 (although not all characters that are convex on both \mathcal{T}_1 and \mathcal{T}_2 necessarily correspond to agreement forests³). Hence there are at most $O(\phi^{2n})$ agreement forests, which is $O(2.619^n)$. It is well-known that d_{TBR} is equal to the number of components in a maximum agreement forest, minus 1 [1]. Hence, again by leveraging Corollary 5 we obtain:

Theorem 7. *Given two unrooted binary trees $\mathcal{T}_1, \mathcal{T}_2$ on the same set of taxa X , where $|X| = n$, $d_{TBR}(\mathcal{T}_1, \mathcal{T}_2) = d_{MAF}(\mathcal{T}_1, \mathcal{T}_2) - 1$ can be computed in time $\Theta(2.619^n \cdot \text{poly}(n))$. Moreover, all agreement forests of \mathcal{T}_1 and \mathcal{T}_2 can be listed in the same time bound.*

6. CONCLUSION

A number of interesting open problems remain. For example, can we characterize non-isomorphic trees that have the same g -spectrum? For a given $k \geq 3$ and n , can we give analytical lower and upper bounds on $g_k(\mathcal{T})$, ranging over the space of all trees \mathcal{T} on n taxa?

7. ACKNOWLEDGEMENTS

We thank Mike Steel for helpful discussions.

REFERENCES

- [1] B. Allen and M. Steel. Subtree transfer operations and their induced metrics on evolutionary trees. *Annals of Combinatorics*, 5:1–15, 2001.
- [2] E. Bachoore and H. Bodlaender. Convex recoloring of leaf-colored trees. *Utrecht University technical report*, 2006.
- [3] M. Bordewich, K. Huber, and C. Semple. Identifying phylogenetic trees. *Discrete Mathematics*, 300(1):30–43, 2005.
- [4] J. Chen, J-H. Fan, and S-H. Sze. Parameterized and approximation algorithms for maximum agreement forest in multifurcating trees. *Theoretical Computer Science*, 562:496–512, 2015.
- [5] A. Dress, K. Huber, J. Koolen, V. Moulton, and A. Spillner. *Basic phylogenetic combinatorics*. Cambridge University Press, 2012.
- [6] M. Fischer and S. Kelk. On the maximum parsimony distance between phylogenetic trees. *Annals of Combinatorics*, 20(1):87–113, 2016.
- [7] W. Fitch. Toward defining the course of evolution: minimum change for a specific tree topology. *Systematic Zoology*, 20(4):406–416, 1971.
- [8] J. Hartigan. Minimum mutation fits to a given tree. *Biometrics*, pages 53–65, 1973.
- [9] S. Kelk and M. Fischer. On the complexity of computing MP distance between binary phylogenetic trees. *Annals of Combinatorics*, 2016. To appear, preprint arXiv:1412.4076.
- [10] S. Kelk, M. Fischer, V. Moulton, and T. Wu. Reduction rules for the maximum parsimony distance on phylogenetic trees. *Theoretical Computer Science*, 2016. To appear, preprint arXiv:1512.07459.
- [11] T. Koshy. *Fibonacci and Lucas numbers with applications*. Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts. Wiley, 2011.
- [12] C. Semple and M. Steel. *Phylogenetics*. Oxford University Press, 2003.
- [13] M. Steel. The complexity of reconstructing trees from qualitative characters and subtrees. *Journal of Classification*, 9(1):91–116, 1992.
- [14] M. Steel and Y. Fu. Classifying and counting linear phylogenetic invariants for the jukes–cantor model. *Journal of Computational Biology*, 2(1):39–47, 1995.

DEPARTMENT OF DATA SCIENCE AND KNOWLEDGE ENGINEERING (DKE), MAASTRICHT UNIVERSITY, P.O. BOX 616, 6200 MD MAASTRICHT, THE NETHERLANDS
E-mail address: `steven.kelk@maastrichtuniversity.nl`
E-mail address: `georgios.stamoulis@maastrichtuniversity.nl`

³However, if $\mathcal{T}_1 = \mathcal{T}_2$ then convex characters and agreement forests are related one-to-one.